

IN THE CLAIMS:

1 1. (CURRENTLY AMENDED) A method for enabling parity declustering in a balanced
2 parity array of a storage system, where an operating system performs the method com-
3 prising the steps of:

4 combining a plurality of unbalanced stripe arrays, each unbalanced stripe array
5 storing an unequal number of blocks per disk, to form the balanced array, the balanced
6 array storing substantially the same number of blocks on all disks, each unbalanced stripe
7 array having parity blocks on a set of storage devices that are disjoint from a set of stor-
8 age devices storing data blocks; and

9 distributing assignment of storage devices to parity groups throughout the bal-
10 anced array.

1 2. (ORIGINAL) The method of Claim 1 further comprising the step of, after a single or
2 double storage device failure, ensuring that all surviving data storage devices are loaded
3 uniformly during reconstruction of the failed storage device or devices.

1 3. (ORIGINAL) The method of Claim 1 wherein the storage system is a filer.

1 4. (ORIGINAL) The method of Claim 1 further comprising the steps of:
2 dividing each storage device into blocks; and
3 organizing the blocks into stripes across the devices, wherein each stripe contains
4 data and parity blocks from each of the devices of the balanced array.

- 1 5. (ORIGINAL) The method of Claim 4 wherein the step of distributing comprises the
2 step of selecting patterns of characters representing data storage devices of a stripe to
3 thereby change the association of the data storage devices with parity groups from stripe
4 to stripe of the balanced array.
- 1 6. (ORIGINAL) The method of Claim 5 wherein the characters are binary numbers.
- 1 7. (ORIGINAL) The method of Claim 5 wherein the characters are ternary numbers.
- 1 8. (ORIGINAL) The method of Claim 1 further comprising the steps of:
 - 2 configuring the balanced array as a RAID-4 style array;
 - 3 initially under-populating the array with storage devices; and
 - 4 adding storage devices until a fully populated array of predetermined size is
5 achieved.
- 1 9. (ORIGINAL) The method of Claim 8 wherein the storage devices are disks.
- 1 10. (CURRENTLY AMENDED) A system that enables parity declustering in a bal-
2 anced parity array of a storage system, the system comprising:
 - 3 a plurality of storage devices, each storage device divided into blocks that are fur-
4 ther organized into stripes, wherein each stripe contains data and parity blocks from each
5 of the devices of the balanced array;
 - 6 a storage operating system including a storage layer configured to implement a
7 parity assignment technique that distributes assignment of devices to parity groups

8 throughout the balanced array such that all storage devices contain the same amount of
9 ~~data or parity information~~; and

10 a processing element configured to execute the operating system to thereby in-
11 voice storage access operations to and from the balanced array in accordance with ~~the a~~
12 concentrated parity technique.

1 11. (ORIGINAL) The system of Claim 10 wherein the storage layer further combines a
2 plurality of unbalanced stripe arrays to form the balanced array, each unbalanced stripe
3 array having parity blocks on a set of storage devices that are disjoint from a set of stor-
4 age devices storing data blocks.

1 12. (ORIGINAL) The system of Claim 11 wherein the storage devices are disks and
2 wherein the storage layer is a RAID layer.

1 13. (ORIGINAL) The system of Claim 12 wherein the RAID layer is implemented in
2 logic circuitry.

1 14. (ORIGINAL) The system of Claim 10 wherein the storage system is a network-
2 attached storage appliance.

1 15. (ORIGINAL) The system of Claim 10 wherein the storage devices are one of video
2 tape, optical, DVD, magnetic tape and bubble memory devices.

1 16. (ORIGINAL) The system of Claim 10 wherein the storage devices are media
2 adapted to store information contained within the data and parity blocks.

1 17. (CURRENTLY AMENDED) Apparatus for enabling parity declustering in a bal-
2 anced parity array of a storage system, the apparatus comprising:

3 means for combining a plurality of unbalanced stripe arrays, each unbalanced
4 stripe array storing an unequal number of blocks per disk, to form the balanced array, the
5 balanced array storing substantially the same number of blocks on all disks, each unbal-
6 anced stripe array having parity blocks on a set of storage devices that are disjoint from a
7 set of storage devices storing data blocks; and

8 means for distributing assignment of devices to parity groups throughout the bal-
9 ~~anced array such that all storage devices contain the same amount of data or parity infor-~~
10 ~~mation.~~

1 18. (ORIGINAL) The apparatus of Claim 17 further comprising:

2 means for dividing each storage device into blocks; and

3 means for organizing the blocks into stripes across the devices, wherein each
4 stripe contains data and parity blocks from each of the devices of the balanced array.

1 19. (ORIGINAL) The apparatus of Claim 18 wherein the means for distributing com-
2 prises means for selecting patterns of characters representing data storage devices of a
3 stripe to thereby change the association of the data storage devices with parity groups
4 from stripe to stripe of the balanced array.

1 20. (CURRENTLY AMENDED) A computer readable medium containing executable
2 program instructions for enabling parity declustering in a balanced parity array of a stor-
3 age system, the executable program instructions comprising program instructions for:

4 combining a plurality of unbalanced stripe arrays, each unbalanced stripe array
5 storing an unequal number of blocks per disk, to form the balanced array, the balanced
6 array storing substantially the same number of blocks on all disks, each unbalanced stripe
7 array having parity blocks on a set of storage devices that are disjoint from a set of stor-
8 age devices storing data blocks; and

9 distributing assignment of devices to parity groups throughout the balanced array
10 such that all storage devices contain the same amount of data or parity information.

1 21. (ORIGINAL) The computer readable medium of Claim 20 further comprising pro-
2 gram instructions for:

3 dividing each storage device into blocks; and

4 organizing the blocks into stripes across the devices, wherein each stripe contains
5 data and parity blocks from each of the devices of the balanced array.

1 22. (ORIGINAL) The computer readable medium of Claim 21 wherein the program in-
2 structions for distributing comprises program instructions for selecting patterns of charac-
3 ters representing data storage devices of a stripe to thereby change the association of the
4 data storage devices with parity groups from stripe to stripe of the balanced array.

1 23. –37. (CANCELLED)

1 38. (PREVIOUSLY PRESENTED) A method for declustering a parity array having a
2 plurality of storage devices, where an operating system performs the method comprising
3 the steps of:

4 assigning a first plurality of data and parity blocks to a first parity group; and

5 assigning a second plurality of data and parity blocks to a second parity group, the
6 first and second parity groups being independent from each other and distributed
7 throughout the plurality of storage devices of the parity array.

1 39. (CANCELLED)

1 40. (PREVIOUSLY PRESENTED) A method for declustering a parity array having a
2 plurality of storage devices, where an operating system performs the method comprising
3 the step of:

4 assigning a plurality of data and parity blocks to a plurality of parity groups, the
5 plurality of parity groups being independent from each other and distributed throughout
6 the plurality of storage devices of the parity array.

1 41. (PREVIOUSLY PRESENTED) A disk array having a declustered parity array,
2 comprising:

3 a plurality of storage devices having a first and second parity group;
4 a first plurality of data and parity blocks assigned to the first parity group; and
5 a second plurality of data and parity blocks assigned to the second parity group,
6 the first and second parity groups being independent from each other and distributed
7 throughout the plurality of storage devices of the parity array.

1 42. (CANCELLED)

1 43. (PREVIOUSLY PRESENTED) A disk array having a declustered parity array,
2 comprising:

3 a plurality of storage devices having a plurality of parity groups; and
4 a plurality of data and parity blocks assigned to the plurality of parity groups, the
5 plurality of parity groups being independent from each other and distributed throughout
6 the plurality of storage devices of the parity array.

1 44. (PREVIOUSLY PRESENTED) A disk array having a declustered parity array,
2 comprising:

3 a plurality of storage devices;
4 means for assigning a first plurality of data and parity blocks to a first parity
5 group; and
6 means for assigning a second plurality of data and parity blocks to a second parity
7 group, the first and second parity groups being independent from each other and distrib-
8 uted throughout the plurality of storage devices of the parity array.

1 45. (CANCELLED)

1 46. (PREVIOUSLY PRESENTED) A disk array having a declustered parity array,
2 comprising:

3 a plurality of storage devices; and
4 means for assigning a plurality of data and parity blocks to a plurality of parity
5 groups, the plurality of parity groups being independent from each other and distributed
6 throughout the plurality of storage devices of the parity array.

1 47. – 54. (CANCELLED)

- 1 55. (NEW) A computer implemented method for enabling parity declustering of a stor-
- 2 age system, the method comprising the steps of:
 - 3 providing a first array of storage devices for storing data blocks and parity blocks,
 - 4 the data blocks organized into at least one parity group associated with the parity blocks,
 - 5 the first array storing an unequal number of blocks on differing ones of the storage de-
 - 6 vices;
 - 7 providing a second array of storage devices for storing data blocks and parity
 - 8 blocks, the data blocks organized into at least one parity group associated with the parity
 - 9 blocks, the second array storing an unequal number of blocks on differing ones of the
 - 10 storage devices;
 - 11 combining the first and second arrays to form a combined array having substan-
 - 12 tially the same number of blocks stored on each storage device of the combined array;
 - 13 and
 - 14 changing the association of data blocks with parity groups in the first array and
 - 15 the second array so that each parity group is associated with data blocks that are distrib-
 - 16 uted substantially uniformly throughout the storage devices that store data blocks in the
 - 17 combined array.

- 1 56. (NEW) The method of claim 55 further comprising the steps of:
 - 2 organizing the data and parity blocks into stripes across the storage devices.
- 1 57. (NEW) The method of claim 56 wherein the step of redistributing comprises the step
- 2 of changing the association of the data storage devices with parity groups from stripe to
- 3 stripe in the combined array.

- 1 58. (NEW) The method of claim 57 wherein the step of changing further comprises the
- 2 step of selecting differing patterns of characters representing data storage devices of a
- 3 stripe.
- 1 59. (NEW) The method of claim 58 wherein the characters are binary numbers.
- 1 60. (NEW) The method of claim 58 wherein the characters are ternary numbers.
- 1 61. (NEW) The method of claim 55 wherein the storage devices are disk drives.
- 1 62. (NEW) A computer implemented method for enabling parity declustering of a stor-
- 2 age array having a plurality of storage devices, the method comprising the steps of:
- 3 dividing each storage device into blocks;
- 4 organizing the blocks into a plurality of stripes across the storage devices,
- 5 wherein each stripe contains data and parity blocks;
- 6 storing data in data blocks and parity information in parity blocks, the parity
- 7 blocks storing parity information for a plurality of parity groups; and
- 8 varying the association of the storage devices to parity groups from stripe to stripe
- 9 in the storage array such that each parity group is associated with data blocks that are dis-
- 10 tributed substantially uniformly throughout the storage devices that store data blocks in
- 11 the storage array.
- 1 63. (NEW) The method of claim 62 wherein the step of changing comprises the step of
- 2 selecting differing patterns of characters representing data storage devices of a stripe to
- 3 thereby change the association of data blocks with parity groups from stripe to stripe of
- 4 the storage array.

1 64. (NEW) An apparatus for enabling parity declustering of a storage system, the appa-
2 ratus comprising:

3 a first array of storage devices for storing data blocks and parity blocks, the data
4 blocks organized into at least one parity group associated with the parity blocks, the first
5 array storing an unequal number of blocks on differing ones of the storage devices;

6 a second array of storage devices for storing data blocks and parity blocks, the
7 data blocks organized into at least one parity group associated with the parity blocks, the
8 second array storing an unequal number of blocks on differing ones of the storage de-
9 vices;

10 a storage operating system configured to combine the first and second arrays to
11 form a combined array having substantially the same number of blocks stored on each
12 storage device of the combined array, and configured to change the association of data
13 blocks with parity groups in the first array and the second array so that each parity group
14 is associated with data blocks that are distributed substantially uniformly = throughout
15 the storage devices that store data blocks in the combined array.

1 65. (NEW) The apparatus of claim 64 wherein each the blocks are organized into stripes
2 across the storage devices.

1 66. (NEW) The apparatus of claim 65 wherein the storage devices are disk drives.

1 67. (NEW) The apparatus of claim 64 wherein the storage devices are one of video tape,
2 optical, DVD, magnetic tape and bubble memory devices.

1 68. (NEW) The apparatus of claim 64 wherein the system is a network-attached storage
2 appliance.

1 69. (NEW) An apparatus for enabling parity declustering of a storage array having a
2 plurality of storage devices, the system comprising:

3 a storage operating system configured to divide each storage device into blocks
4 and organize the blocks into a plurality of stripes across the storage devices, wherein each
5 stripe contains data and parity blocks and store data in data blocks and parity information
6 in parity blocks, the parity blocks storing parity information for a plurality of parity
7 groups;

8 the storage operating system further configured to vary the association of the stor-
9 age devices to parity groups from stripe to stripe in the storage array such that, each par-
10 ity group is associated with data blocks that are distributed substantially uniformly
11 throughout the storage devices that store data blocks in the storage array.

1 70. (NEW) The system of claim 68 wherein the storage operating system is configured
2 to select differing patterns of characters representing data storage devices of a stripe to
3 thereby change the association of data blocks with parity groups from stripe to stripe of
4 the storage array.

1 71. (NEW) A system that enables parity declustering of a storage system, the system
2 comprising:

3 means for providing a first array of storage devices for storing data blocks and
4 parity blocks, the data blocks organized into at least one parity group associated with the

5 parity blocks, the first array storing an unequal number of blocks on differing ones of the
6 storage devices;

7 means for providing a second array of storage devices for storing data blocks and
8 parity blocks, the data blocks organized into at least one parity group associated with the
9 parity blocks, the second array storing an unequal number of blocks on differing ones of
10 the storage devices;

11 means for combining the first and second arrays to form a combined array having
12 substantially the same number of blocks stored on each storage device of the combined
13 array; and

14 means for changing the association of data blocks with parity groups in the first
15 array and the second array so that each parity group is associated with data blocks that are
16 distributed substantially uniformly throughout the storage devices that store data blocks
17 in the combined array.

1 72. (NEW) An apparatus for enabling parity declustering of a storage array having a
2 plurality of storage devices, the system comprising:

3 means for dividing each storage device into blocks;

4 means for organizing the blocks into a plurality of stripes across the storage de-
5 vices, wherein each stripe contains data and parity blocks;

6 means for storing data in data blocks and parity information in parity blocks, the
7 parity blocks storing parity information for a plurality of parity groups; and

8 means for varying the association of the storage devices to parity groups from
9 stripe to stripe in the storage array such that, each parity group is associated with data
10 blocks that are distributed substantially uniformly throughout the storage devices that
11 store data blocks in the storage array.

1 73. (NEW) A computer readable medium containing program instructions for execution
2 on a processor, the executable program instructions comprising program instructions for:

3 dividing each storage device into blocks;

4 organizing the blocks into a plurality of stripes across the storage devices,
5 wherein each stripe contains data and parity blocks;

6 storing data in data blocks and parity information in parity blocks, the parity
7 blocks storing parity information for a plurality of parity groups; and

8 varying the association of the storage devices to parity groups from stripe to stripe
9 in the storage array such that, each parity group is associated with data blocks that are
10 distributed substantially uniformly throughout the storage devices that store data blocks
11 in the storage array.

1 74. (NEW) The apparatus of claim 17 wherein each unbalanced stripe array has fewer
2 parity blocks per disk than data blocks per disk.

1 75. (NEW) The computer readable medium of claim 20 wherein each unbalanced stripe
2 array has fewer parity blocks per disk than data blocks per disk.

1 76. (NEW) An apparatus for parity declustering in a storage system, the apparatus com-
2 prising:

3 a storage operating system configured to combine a plurality of first arrays of
4 storage devices, each first array storing an unequal number of blocks per storage device,
5 to form a second array, the second array storing substantially the same number of blocks
6 on all storage devices;

7 the storage operating system further configured to redistribute the assignment
8 storage devices to parity groups in the second array so that each storage device will have
9 a substantially equal number of blocks associated with each parity group.

1 77. (NEW) An method for parity declustering in a storage system, the method comprising
2 the steps of:

3 combining a plurality of first arrays of storage devices, each first array storing an
4 unequal number of blocks per storage device, to form a second array, the second array
5 storing substantially the same number of blocks on all storage devices;

6 the storage operating system further configured to redistribute the assignment
7 storage devices to parity groups in the second array so that each storage device will have
8 a substantially equal number of blocks associated with each parity group.

1 78. (NEW) A computer implemented method for enabling parity declustering of a stor-
2 age system, the method comprising the steps of:

3 providing a first array of storage devices for storing data blocks and parity blocks,
4 the data blocks organized into at least one parity group associated with the parity blocks,
5 the first array storing an unequal number of blocks on differing ones of the storage de-
6 vices;

7 providing a second array of storage devices for storing data blocks and parity
8 blocks, the data blocks organized into at least one parity group associated with the parity

9 blocks, the second array storing an unequal number of blocks on differing ones of the
10 storage devices;

11 combining the first and second arrays to form a combined array having substan-
12 tially the same number of blocks stored on each storage device of the combined array.

1 79. (NEW) An apparatus for enabling parity declustering of a storage system, the appa-
2 ratus comprising:

3 a first array of storage devices for storing data blocks and parity blocks, the data
4 blocks organized into at least one parity group associated with the parity blocks, the first
5 array storing an unequal number of blocks on differing ones of the storage devices;

6 a second array of storage devices for storing data blocks and parity blocks, the
7 data blocks organized into at least one parity group associated with the parity blocks, the
8 second array storing an unequal number of blocks on differing ones of the storage de-
9 vices;

10 a storage operating system configured to combine the first and second arrays to
11 form a combined array having substantially the same number of blocks stored on each
12 storage device of the combined array.